

REMARKS

A. BACKGROUND OF THIS RESPONSE AND OVERVIEW

This paper is being filed to initiate an interference with Leighton *et al.* U.S. Patent No. 6,108,703.

The origin of newly presented claims 41-61 is summarized on the following TABLE, for the Examiner's ease of cross-reference:

NEW CLAIM	LEIGHTON '703 CLAIM
41	1
42	4
43	5
44	6
45	9
46	10
47	11
48	14
49	15
50	16
51	17
52	18
53	19
54	20
55	21
56	22
57	23
58	31
59	32
60	33
61	34

Following entry of the above claim amendments, claims 25-61 will be pending in this application.

**B. APPLICATION OF COPIED CLAIMS TO SPECIFICATION
AND PRIORITY APPLICATION DISCLOSURE**

Pursuant to 37 CFR § 1.607(a)(5), the table of **APPENDIX I** hereto identifies the supporting disclosure in the present application for new claims 41-61, which have been copied in identical form from the Leighton '703 Patent.

The content of each column of the table of **APPENDIX I** is as follows:

- The first column shows the claim number in the present application and the number of the corresponding claim in the Leighton '703 Patent, the claim text as it appears in the Leighton '703 Patent.
- The second column recites page and line numbers where support for each copied claim is found in the present application.

It is clear from **APPENDIX I** that each copied claim finds support back to Appln.

No. 09/021,506 filed February 10, 1998, the priority date for the present application.

**C. IDENTIFICATION OF INTERFERING PATENT AND
EFFECTIVE FILING DATES**

Pursuant to 37 CFR § 1.607(a)(1), applicants seek to provoke an interference with Leighton *et al.* U.S. Patent No. 6,108,703, filed May 19, 1999, issued August 22, 2000. Leighton's priority is claimed from Provisional U.S. Application No. 60/092,710, filed July 14, 1998.

D. PROPOSED COUNT AND CORRESPONDING CLAIMS

Pursuant to 37 CFR § 1.607(a)(2), applicants propose the following Counts.

Counts 1 and 2 correspond to claims 1 and 14 of the Leighton '703 patent, respectively. Count 25 corresponds to pending claim 25 of the present application.

Count 1:

A distributed hosting framework operative in a computer network in which users of client machines connect to a content provider server, the framework comprising:

a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path;

a set of content servers, distinct from the content provider server, for hosting at least some of the embedded objects of web pages that are normally hosted by the content provider server;

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level domain name service (DNS) resolution;

wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the content provider server and the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers.

Count 2:

A method of serving a page supported at a content provider server, the page comprising a markup language base document having associated therewith a set of embedded objects, each embedded object identified by a URL, comprising the steps of:

rewriting the URL of an embedded object to generate a modified URL, the modified URL including a new hostname prepended to an original hostname, wherein the original hostname is maintained as part of the modified URL for use in retrieving the embedded object whenever a cached copy of the embedded object is not available;

in response to a request to serve the page received at the content provider site, serving the page with the modified URL;

attempting to serve the embedded object from a content server other than the content provider server as identified by the new hostname; and

if the cached copy of the embedded object is not available from the content server, serving the embedded object from the content provider server.

Count 3

A method comprising:

obtaining a first resource containing a reference to a second resource; and

replacing, in the first resource, the reference to the second resource with a different resource reference.

E. IDENTIFICATION OF CLAIMS CORRESPONDING TO THE COUNTS

Pursuant to 37 CFR § 1.607(a)(3) and (4), the following claims of the Leighton '703 Patent and of the present application are identified as corresponding to the counts:

a) Claims Of the Leighton '703 Patent:

Count 1

- Claim 1 of the Leighton '703 Patent corresponds exactly to Count 1.

APPLICATION of FARBER et al. – Appln. No. 09/612,598

- At least claims 19 and 34 correspond substantially to Count 1.

Count 2

- Claim 14 of the Leighton '703 Patent corresponds exactly to Count 2.
- At least claims 15, 17, 23 correspond substantially to Count 2.

Count 3

- All claims of the Leighton '703 Patent correspond to Count 3.

b) Claims Of Present Application

Count 1

Copied claim 41 of the present application (claim 1 of the Leighton '703 Patent) corresponds exactly to Count 1. Copied claims 53 and 61 correspond exactly to Count 1.

Count 2

Copied claim 48 of the present application corresponds exactly to count 2.
At least copied claims 49, 51, and 57 correspond substantially to Count 2.

Count 3

Claim 25 of the present application corresponds exactly to count 3.

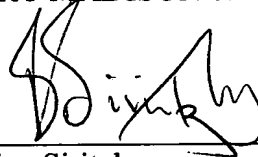
F. CONCLUSION

All applicable requirements of 37 CFR § 1.607 having been complied with, it is respectfully requested that an interference be declared between the present application and with Leighton *et al.*, U.S. Patent No. 6,108,703.

Respectfully submitted,

PILLSBURY MADISON & SUTRO, LLP

By



Brian Siritzky
Reg. No. 37497
Tel. No.: (202) 861-3702
Fax No.: (202) 822-0944

1100 New York Avenue, N.W.
Ninth Floor
Washington, D.C. 20005-3918
(202) 861-3000
30096951v1

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
41/1. A distributed hosting framework operative in a computer network in which users of client machines connect to a content provider server, the framework comprising:	See generally Fig. 1 (e.g., client 106 and origin server 102) and corresponding description, e.g., at pgs. 6 to 8.
a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path;	See generally Fig. 3, “ B5 Rewrite Resource ” and corresponding description, e.g., at pgs. 15, lines 14-15, (“If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it”).
	Resource rewriting is explained and described in greater detail in the section titled “ Rewriting HTML Resources ” on pgs. 30 and 31.
	One form of the modified URL is specified at pg. 14, lines 13-16, which states:
	<p>D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is “http”, the preferred embodiment uses the following format:</p> <p style="padding-left: 40px;"><i>http://<repeater>/<server>/<path></i></p> <p>See also, e.g., “create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information.” <i>Pg. 14, lines 8-10.</i></p>
a set of content servers, distinct from the content provider server, for hosting at least some of the embedded objects of web pages that are normally hosted by the content provider server;	See generally Fig. 1, repeaters 104. Each repeater 104a, 104b, and 104c replicates some or all of the information available on the origin server 102 as well as information available on other origin servers in the network 100.” <i>Pg. 6, lines 16-18.</i>
at least one first level name server that provides a first level domain name service (DNS) resolution; and	“The browser extracts the host ... name ... and uses a domain name server (DNS) to look up the network ... address of the corresponding server.” <i>Pg. 11, lines 13-15.</i>
at least one second level name server that provides a second level domain name service (DNS) resolution;	The reflector 108 intercepts the request for the origin server 102 and determines where the request should be served. I.e., it determines the address of a server to handle the request. E.g., <i>Pg. 8, lines 2-4.</i>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
wherein in response to requests for the web page, generated by the client machines	In addition, it is well known in the art to use multiple levels of DNS. E.g., root name servers. See, e.g., Albitz, P. and Liu, C. "DNS and BIND in a Nutshell", O'Reilly & Assoc., 1992, Chap. 2, "How Does DNS Work?", pgs. 13-38, especially pgs. 29-30. ¹ See generally Fig. 2, "A5 Receive Request" and corresponding description. "The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." Pg. 19, lines 9-11.
the web page including the modified embedded object URL is served from the content provider server and the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers.	See generally Fig. 2, "A7 Send Reply with Resource" and corresponding description. "When a browser receives a REDIRECT response (as produced in B3), it reissues a request for the resource using the new resource identifier (URL) (A1-A5). Because the new identifier refers to a repeater instead of the origin server, the browser now sends a request for the resource to the repeater." Pg. 17, line 27 to pg. 18, line 3.
Claim (Appln. Claim No./Patent Claim No.) 42/4. The hosting framework as described in claim 1 wherein a given one of the set of servers includes a buddy server for assuming the hosting responsibilities of the given one of the set of servers upon a given failure condition.	Support in Application See generally Fig. 1, repeaters 104. "Each repeater 104a, 104b, and 104c replicates some or all of the information available on the origin server 102 as well as information available on other origin servers in the network 100." Pg. 6, lines 16-18. See also section titled "Repeater Network Resilience" at pg. 39, line 3 to pg. 41, line 4. "If a master fails ... another repeater will take over the role of master" Pg. 39, lines 9-10.

¹ A courtesy copy of this Chapter is provided herewith for the Examiner.

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
43/5. The hosting framework as described in claim 1 wherein the second level name server includes a load balancing mechanism that balances loads across a subset of the set of servers.	Reflector 108 includes mechanism to determine best repeater to reflect request. E.g., Pg. 14, lines 1-2. “An appropriate repeater is one which is not too heavily loaded.” Pg. 19, lines 24-25.
Claim (Appln. Claim No./Patent Claim No.)	Support in Application
43/6. The hosting framework as described in claim 5 wherein the load balancing mechanism minimizes the amount of replication required for the embedded objects while not exceeding a capacity of any of the set of servers.	See generally Selecting Best Repeater , pg. 19, line 19 to pg. 23, line 27.. Also, e.g., “If a repeater’s load exceeds its configured capacity, an alarm message is sent to the repeater network administrator.” Pg. 21, lines 17-19.
Claim (Appln. Claim No./Patent Claim No.)	Support in Application
45/9. The hosting framework as described in claim 1 wherein the first level name server includes a network map for use in directing a request for the embedded object generated by a client.	<p>“The BRS relies on three pre-computed tables, namely the Group Reduction Table, the Link Cost Table, and the Load Table. These three tables (described below) are computed off-line and downloaded to each reflector by its contact in the repeater network” pg. 20, lines 4-7.</p> <p>“The Group Reduction Table and Link Cost Table used in BRS processing are created and regularly updated by an independent procedure referred to herein as <i>NetMap</i>.” Pg. 24, lines 2-5.</p> <p>“The <i>NetMap</i> procedure ... [creates a] database ... [that] contains essential information used for further processing, namely (1) the identity of each group, (2) the set of IP addresses in a given group, (3) the presence of links between groups indicating paths over which information may travel, and (4) the cost of sending data over a given link.” Pg. 24, line 20 to pg. 25, line 1.</p>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
46/10. The hosting framework as described in claim 1 wherein a server in the set of servers includes a gating mechanism for maintaining overall traffic for a given embedded object within specified limits.	<p>"... the reflector 108 can be configured so that requests from certain network addresses ... are never reflected. Also, the reflector may choose not to reflect requests because the reflector is exceeding its committed aggregate information rate." Pg. 16, lines 11-14.</p> <p>See also the Section titled "Enforcing the Committed Aggregate Information Rate" at pgs. 37 to 39.</p>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
47/11. The hosting framework as described in claim 10 wherein the gating mechanism comprises: means for determining whether a number of requests for the given embedded object exceeds a given threshold; and means responsive to the determining means for restricting service of the given embedded object.	<p>"a maximum capacity setting is set. The maximum capacity indicates the point at which the repeater is considered to be fully loaded" Pg. 21, lines 10-11.</p> <p>See generally the section titled "Enforcing Committed Aggregate Information Rate" starting at Pg. 37.</p> <p>"The repeater network monitors and limits the aggregate rate at which data is served on behalf of a given subscriber by all repeaters" pg. 37, lines 3-4.</p> <p>"For each subscriber, a "threshold aggregate information rate" (TAIR) is configured and maintained at the master repeater. ... When a reflector receives a request, it determines whether its most recently calculated MAIR is greater than its TAIR. If this is the case, the reflector probabilistically decides whether to suppress reflection" pg. 37, lines 10-24.</p>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
48/14. A method of serving a page supported at a content provider server,	See generally Fig. 1 and its corresponding description.
the page comprising a markup language base document having associated therewith a set of embedded objects, each embedded object identified by a URL, comprising the steps of:	“In the case of the Internet in general and the World Wide Web specifically, documents can be created using a standardized form known as the Hypertext Markup Language (HTML). In HTML, a document consists of data (text, images, sounds, and the like), including links to other sections of the same document or to other documents. The links are generally provided as URLs, and can be in relative or absolute form.” Pg. 9, lines 10-16.
rewriting the URL of an embedded object to generate a modified URL, the modified URL including a new hostname prepended to an original hostname,	See generally the section titled “ Rewriting HTML Resources ” at pgs. 30 and 31. All, generally Fig. 3, “B5 Rewrite Resource” and corresponding description, e.g., at pgs. 15, lines 14-15, (“If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it”). One form of the modified URL is specified at pg. 14, lines 13-16, which states: D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is “http”, the preferred embodiment uses the following format: <i>http://<repeater>/<server>/<path></i> See also, e.g., “create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information.” Pg. 14, lines 8-10.
wherein the original hostname is maintained as part of the modified URL for use in retrieving the embedded object whenever a cached copy of the embedded object is not available;	“The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client.” Pg. 19, lines 9-11, emphasis added.
in response to a request to serve the page received at the content provider site, serving the page with the modified URL;	“When a browser receives a REDIRECT response (as produced in B3), it reissues a request for the resource using the new resource identifier (URL) (A1-A5). Because the new identifier refers to a repeater instead of the origin server, the browser now sends a request for the resource to the repeater” Pg. 17, line 27 to pg. 18, line 3.
attempting to serve the embedded object from	“The repeater then determines whether the requested resource is cached locally. If the

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
a content server other than the content <u>provider</u> server as identified by the new hostname; and	requested resource is in the <u>repeater's</u> cache it is retrieved. On the other hand, if a valid copy of the requested resource is not in the <u>repeater's</u> cache, the repeater modifies the incoming URL, creating a request that it issues directly to the originating <u>reflector</u> which processes it (as in B1-B6)." Pg. 18, line 20 to pg. 19, line 7.
if the cached copy of the embedded object is not available from the <u>content server</u> , serving the embedded object from the <u>content provider</u> server.	"If a resource is not cached locally, the cache can query its "peer caches" to see if one of them contains the resource, before or at the same time as requesting the resource from the <u>reflector/origin server</u> ." Pg. 19, lines 2-8.

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
49/15. A method of serving a page and an associated page object, wherein the page is stored on a content provider server and copies of the page object are stored on a set of content servers distinct from the content provider server, comprising the steps of:	See generally Fig. 1, repeaters 104. "Each repeater 104a, 104b, and 104c replicates some or all of the information available on the origin server 102 as well as information available on other origin servers in the network 100." Pg. 6, lines 16-18. Also, "A hypertext document may contain any number of links to other documents, and each of those other documents may be on a different server in a different part of the world." pg. 9, lines 20-23.
(a) modifying a URL for the page object to include a hostname prepended to a content provider-supplied domain name and path;	See generally the section titled "Rewriting HTML Resources" at pgs.. 30 and 31. All, generally Fig. 3, "B5 Rewrite Resource" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it"). One form of the modified URL is specified at pg. 14, lines 13-16, which states: D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: <code>http://<repeater>/<server>/<path></code>
	See also, e.g., "create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
(b) serving the page from the content provider server with the modified URL;	provide this information.” Pg. 14, lines 8-10. Fig. 5, “Send GET Request to Repeater” and corresponding description. “C4. The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client.” Pg. 19, lines 9-11.
(c) responsive to a browser query to resolve the hostname, identifying a given one of the set of content servers from which the object may be retrieved; and	“The selection by the reflector of an appropriate repeater to handle the request can be done in a number of ways. In the preferred embodiment, it is done by first pre-partitioning the network into “cost groups” and then determining which cost group the client is in. Next, from a plurality of repeaters in the network, a set of repeaters is selected, the members of the set having a low cost relative to the cost group which the client is in. ... Then one member of the set is selected ... as the best repeater.” Pg. 4, line 27 to pg. 5, line 6. See also Fig. 6, “E3 Select small set of repeaters” and “E4 Select Element from set” and corresponding description at pg. 23.
(d) returning to the browser an IP address of the identified content server to enable the browser to attempt to retrieve the object from that content server.	“When a browser receives a REDIRECT response (as produced in B3), it reissues a request for the resource using the new resource identifier (URL) (A1-A5). Because the new identifier refers to a repeater instead of the origin server, the browser now sends a request for the resource to the repeater” Pg. 17, line 27 to pg. 18, line 3.

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
50/16. The method as described in claim 15 wherein the copies of the page object are stored on a subset of the set of content servers.	“Each repeater 104a, 104b, and 104c replicates <u>some or all</u> of the information available on the origin server 102 as well as information available on other origin servers in the network 100.” Pg. 6, lines 16-18, emphasis added.

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
<p>51/17. A content delivery method, comprising: tagging an embedded object in a page to resolve to a domain other than a content provider domain by prepending given data to a content provider-supplied URL to generate an alternate resource locator (ARL);</p>	<p>“A hypertext document may contain any number of links to other documents, and each of those other documents may be on a different server in a different part of the world.” <i>pg. 9, lines 20-23.</i></p> <p>See generally the section titled “Rewriting HTML Resources” at pgs.. 30 and 31. All, generally Fig. 3, “B5 Rewrite Resource” and corresponding description, e.g., at pgs. 15, lines 14-15, (“If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it”). One form of the modified URL is specified at pg. 14, lines 13-16, which states:</p> <p>D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is “http”, the preferred embodiment uses the following format: <i>http://<repeater>/<server>/<path></i></p> <p>See also, e.g., “create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information.” <i>Pg. 14, lines 8-10</i></p>
serving the page from a content provider server with the ARL; and	Fig. 5, “Send GET Request to Repeater” and corresponding description.
resolving the ARL to identify a content server in the domain; and	Fig. 3 “Reply is REDIRECT?” “Go back to A1 with new URL” and corresponding description.
serving the embedded object from the identified content server.	“The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client.” <i>Pg. 19, lines 9-11.</i>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
52/18. The method as described in claim 17 wherein the step of resolving the ARL comprises: utilizing a requesting user's location and data identifying then-current Internet traffic conditions to identify the content server.	"reflector 108 can be configured so that requests from certain network addresses (e.g., requests from clients on the same local area network as the reflector itself) are never reflected. Also, the reflector may choose not to reflect requests because the reflector is exceeding its committed aggregate information rate" <i>Pg. 16, lines 11-14.</i>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
53/19. A content delivery service, comprising: replicating a set of page objects across a wide area network of content servers managed by a domain other than a content provider domain; for a given page normally served from the content provider domain, tagging the embedded objects of the page so that requests for the page objects resolve to the domain instead of the content provider domain;	<p>See generally Fig. 1, repeaters 104. "Each repeater 104a, 104b, and 104c replicates some or all of the information available on the origin server 102 as well as information available on other origin servers in the network 100." <i>Pg. 6, lines 16-18.</i></p> <p>"See generally the section titled "Rewriting HTML Resources" at pgs. 30 and 31. All, generally Fig. 3, "B5 Rewrite Resource" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it").</p> <p>"Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." <i>Pg. 30, lines 14-15.</i> "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." <i>Pg. 31, lines 10-14.</i></p> <p>One form of the modified URL is specified at pg. 14, lines 13-16, which states:</p> <p>D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: <i>http://<repeater>/<server>/<path></i></p>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
	See also, e.g., “create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information.” <i>Pg. 14, lines 8-10</i>
responsive to a request for the given page received at the content provider domain, serving the given page from the content provider domain; and	“The repeater then determines whether the requested resource is cached locally. If the requested resource is in the repeater’s cache it is retrieved. On the other hand, if a valid copy of the requested resource is not in the repeater’s cache, the repeater modifies the incoming URL, creating a request that it issues directly to the originating reflector which processes it (as in B1-B6).” <i>Pg. 18, line 20 to pg. 19, line 7.</i>
serving at least one embedded object of the given page from a given content server in the domain instead of from the content provider domain.	“When a browser receives a REDIRECT response (as produced in B3), it reissues a request for the resource using the new resource identifier (URL) (A1-A5). Because the new identifier refers to a repeater instead of the origin server, the browser now sends a request for the resource to the repeater” <i>Pg. 17, line 27 to pg. 18, line 3.</i>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
54/20. The content delivery method as described in claim 19 wherein the serving step comprises:	
for each embedded object, identifying one or more content servers from which the embedded object may be retrieved.	<p>“If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it” at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16.</p> <p>Resource rewriting is explained and described in greater detail in the section titled “Rewriting HTML Resources” on pgs. 30 and 31.</p> <p>“Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector).” <i>Pg. 30, lines 14-15.</i></p> <p>“For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater.” <i>Pg. 31, lines 10-14.</i></p>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
55/21. The method as described in claim 20 wherein the identifying step comprises: resolving a request to the domain as a function of a requesting user's location.	"An appropriate repeater is one which is not too heavily loaded and which is not too far from the client in terms of some measure of network distance." Pg. 19, line 24 to pg. 20, line 25.

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
56/22. The method as described in claim 21 wherein the identifying step comprises: resolving a request to the domain as a function of a requesting user's location and then-current Internet traffic conditions.	"reflector 108 can be configured so that requests from certain network addresses (e.g., requests from clients on the same local area network as the reflector itself) are never reflected. Also, the reflector may choose not to reflect requests because the reflector is exceeding its committed aggregate information rate" Pg. 16, lines 11-14.

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
57/23. A method for Internet content delivery, comprising: at the content provider server, modifying at least one embedded object URL of a page to include a hostname prepended to a domain name and a path normally used to retrieve the embedded object;	See generally the section titled "Rewriting HTML Resources" at pgs. 30 and 31. Also, generally Fig. 3, "B5 Rewrite Resource" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it"). One form of the modified URL is specified at pg. 14, lines 13-16, which states: D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: <code>http://<repeater>/<server>/<path></code>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
	See also, e.g., "create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information." <i>Pg. 14, lines 8-10.</i>
responsive to a request for the page issued from a client machine, serving the page with the modified embedded object URL to the client machine from the content provider server;	See generally Fig. 3 and corresponding description at pgs. 14-15. "upon receipt of a request, B1 The reflector ... analyzes the request to determine whether or not to reflect the request. ... B5. If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it, The resource, possibly as modified by rewriting, is then returned in a reply to the requesting client" <i>Pg. 13, line 2 to pg. 15, line 18.</i> The embedded object will have a modified URL. <i>Pg. 13, line 2 to pg. 15, line 18.</i>
responsive to a request for the embedded object, resolving the hostname to an IP address of a content server, other than the content provider server, that is likely to host the embedded object; and attempting to serve the embedded object to the client from the content server.	"The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." <i>Pg. 19, lines 9-12.</i>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
58/31. The method as described in claim 23 wherein the page is formatted according to a markup language.	"In the case of the Internet in general and the World Wide Web specifically, documents can be created using a standardized form known as the Hypertext Markup Language (HTML)." <i>Pg. 9, lines 10-12.</i>

Claim (Appln. Claim No./Patent Claim No.)	Support in Application
59/32. The method as described in claim 23 further including the step of rewriting the embedded object URL as the content provider modifies the page.	See generally the section titled "Rewriting HTML Resources" at pgs.. 30 and 31.
Claim (Appln. Claim No./Patent Claim No.)	Support in Application
60/33. The method as described in claim 23 wherein the step of resolving the hostname includes:	
identifying a subset of content servers that may be available to serve the embedded object based on a location of the client machine and current Internet traffic conditions; and	<p>"The selection ... of an appropriate repeater to handle the request can be done in a number of ways. In the preferred embodiment, it is done by first pre-partitioning the network into "cost groups" and then determining which cost group the client is in. Next, from a plurality of repeaters in the network, a set of repeaters is selected, the members of the set having a low cost relative to the cost group which the client is in." Pg. 4, line 27 to pg. 5, line 5.</p> <p>"... requests from certain network addresses (e.g., requests from clients on the same local area network as the reflector itself) are never reflected. Also, the reflector may choose not to reflect requests because the reflector is exceeding its committed aggregate information rate" Pg. 16, lines 11-14</p> <p>See also pg. 15, lines 26-28.</p>
identifying the content server from the subset of content servers.	"Then one member of the set is selected, ..., as the best repeater." Pg. 5, lines 4-5.